Constant-Round Group Key Exchange from the Ring-LWE Assumption

Daniel Apon[1], Dana Dachman-Soled[2], Huijing Gong[2], and Jonathan Katz[2]

[1] National Institute of Standards and Technology, USA
daniel.apon@nist.gov
[2] University of Maryland, College Park, USA
danadach@ece.umd.edu, {gong, jkatz}@cs.umd.edu

Abstract. Group key-exchange protocols allow a set of N parties to agree on a shared, secret key by communicating over a public network. A number of solutions to this problem have been proposed over the years, mostly based on variants of Diffie-Hellman (two-party) key exchange. To the best of our knowledge, however, there has been almost no work looking at candidate *post-quantum* group key-exchange protocols.

Here, we propose a constant-round protocol for unauthenticated group key exchange (i.e., with security against a passive eavesdropper) based on the hardness of the Ring Learning With Errors (Ring-LWE) problem. By applying the Katz-Yung compiler using any post-quantum signature scheme, we obtain a (scalable) protocol for *authenticated* group key exchange with post-quantum security. Our protocol is constructed by generalizing the Burmester-Desmedt protocol to the Ring-LWE setting, which requires addressing several technical challenges.

Keywords: Ring learning with errors, Post-quantum cryptography, Group key exchange

1 Introduction

Protocols for (authenticated) key exchange are among the most fundamental and widely used cryptographic primitives. They allow parties communicating over an insecure public network to establish a common secret key, called a *session key*, permitting the subsequent use of symmetric-key cryptography for encryption and authentication of sensitive data. They can be used to instantiate so-called "secure channels" upon which higher-level cryptographic protocols often depend.

Most work on key exchange, beginning with the classical paper of Diffie and Hellman, has focused on two-party key exchange. However, many works have also explored extensions to the *group* setting [21, 29, 15, 30, 5, 6, 25, 14, 12, 13, 11, 17, 22, 16, 8, 2, 1, 24, 9, 31] in which N parties wish to agree on a common session key that they can each then use for encrypted/authenticated communication with the rest of the group.

The recent effort by the National Institute of Standards and Technology (NIST) to evaluate and standardize one or more quantum-resistant public-key cryptosystems is entirely focused on digital signatures and two-party key encapsulation/key exchange,[1] and there has been an extensive amount of research over the past decade focused on designing such schemes. In contrast, we are aware of almost no[2] work on *group* key-exchange (GKE) protocols with post-quantum security beyond the observation that a post-quantum group key-exchange protocol can be constructed from any post-quantum two-party protocol by having a designated group manager run independent two-party protocols with the $N-1$ other parties, and then send a session key of its choice to the other parties encrypted/authenticated using each of the resulting keys. Such a solution is often considered unacceptable since it is highly asymmetric, requires additional coordination, is not contributory, and puts a heavy load on a single party who becomes a central point of failure.

1.1 Our Contributions

In this work, we propose a constant-round group key-exchange protocol based on the hardness of the Ring-LWE problem [27], and hence with (plausible) post-quantum security. We focus on constructing an *unauthenticated* protocol—i.e., one secure against a passive eavesdropper—since known techniques such as the Katz-Yung compiler [24] can then be applied to obtain an *authenticated* protocol secure against an active attacker.

The starting point for our work is the two-round group key-exchange protocol by Burmester and Desmedt [15, 16, 24], which is based on the decisional Diffie-Hellman assumption. Assume a group $\mathbb{G}$ of prime order q and a generator $g \in \mathbb{G}$ are fixed and public. The Burmester-Desmedt protocol run by parties $P_0, \ldots, P_{N-1}$ then works as follows:

1. In the first round, each party P_i chooses uniform $r_i \in \mathbb{Z}_q$ and broadcasts $z_i = g^{r_i}$ to all other parties.
2. In the second round, each party P_i broadcasts $X_i = (z_{i+1}/z_{i-i})^{r_i}$ (where the parties' indices are taken modulo N).

Each party P_i can then compute its session key sk_i as

$$\mathsf{sk}_i = (z_{i-1})^{Nr_i} \cdot X_i^{N-1} \cdot X_{i+1}^{N-2} \cdots X_{i+N-2}.$$

One can check that all the keys are equal to the same value $g^{r_0 r_1 + \cdots + r_{N-1} r_0}$.

In attempting to adapt their protocol to the Ring-LWE setting, we could fix a ring R_q and a uniform element $a \in R_q$. Then:

1. In the first round, each party P_i chooses "small" secret value $s_i \in R_q$ and "small" noise term $e_i \in R_q$ (with the exact distribution being unimportant in the present discussion), and broadcasts $z_i = as_i + e_i$ to the other parties.

[1] Note that CPA-secure key encapsulation is equivalent to two-round key-exchange (with passive security).

[2] The protocol of Ding et al.[19] has no security proof; the work of Boneh et al.[10] shows a framework for constructing a group key-exchange protocol with plausible post-quantum security but without a concrete instantiation.

2. In the second round, each party P_i chooses a second "small" noise term $e'_i \in R_q$ and broadcasts $X_i = (z_{i+1} - z_{i-i}) \cdot s_i + e'_i$.

Each party can then compute a session key b_i as

$$b_i = N \cdot s_i \cdot z_{i-1} + (N-1) \cdot X_i + (N-2) \cdot X_{i+1} + \cdots + X_{i+N-2}.$$

The problem, of course, is that (due to the noise terms) these session keys computed by the parties will *not* be equal. They will, however, be "close" to each other if the $\{s_i, e_i, e'_i\}$ are all sufficiently small, so we can add an additional reconciliation step to ensure that all parties agree on a common key k.

This gives a protocol that is correct, but proving security (even for a passive eavesdropper) is more difficult than in the case of the Burmester-Desmedt protocol. Here we informally outline the main difficulties and how we address them. First, we note that trying to prove security by direct analogy to the proof of security for the Burmester-Desmedt protocol (cf. [24]) fails; in the latter case, it is possible to use the fact that, for example,

$$(z_2/z_0)^{r_1} = z_1^{r_2 - r_0},$$

whereas in our setting the analogous relation does not hold. In general, the natural proof strategy here is to switch all the $\{z_i\}$ values to uniform elements of R_q, and similarly to switch the $\{X_i\}$ values to uniform subject to the constraint that their sum is approximately 0 (i.e., subject to the constraint that $\sum_i X_i \approx 0$). Unfortunately this cannot be done by simply invoking the Ring-LWE assumption $O(N)$ times; in particular, the first time we try to invoke the assumption, say on the pair $(z_1 = as_1 + e_1,\ X_1 = (z_2 - z_0) \cdot s_1 + e'_1)$, we need $z_2 - z_0$ to be uniform—which, in contrast to the analogous requirement in the Burmester-Desmedt protocol (for the value z_2/z_0), is not the case here. Thus, we must somehow break the circularity in the mutual dependence of the $\{z_i, X_i\}$ values.

Toward this end, let us look more carefully at the distribution of $\sum_i X_i$. We may write

$$\textstyle\sum_i X_i = \sum_i (e_{i+1}s_i - e_{i-1}s_i) + \sum_i e'_i.$$

Consider now changing the way X_0 is chosen: that is, instead of choosing $X_0 = (z_1 - z_{N-1})s_0 + e'_0$ as in the protocol, we instead set $X_0 = -\sum_{i=1}^{N-1} X_i + e'_0$ (where e'_0 is from the same distribution as before). Intuitively, as long as the standard deviation of e'_0 is large enough, these two distributions of X_0 should be "close" (as they both satisfy $\sum_i X_i \approx 0$). This, in particular, means that we need the distribution of e'_0 to be different from the distribution of the $\{e'_i\}_{i>0}$, as the standard deviation of the former needs to be larger than the latter.

We can indeed show that when we choose e'_0 from an appropriate distribution then the Rényi divergence between the two distributions of X_0, above, is bounded by a polynomial. With this switch in the distribution of X_0, we have broken the circularity and can now use the Ring-LWE assumption to switch the distribution of z_0 to uniform, followed by the remaining $\{z_i, X_i\}$ values.

Unfortunately, bounded Rényi divergence does not imply statistical closeness. However, polynomially bounded Rényi divergence *does* imply that any event occurring with negligible probability when X_0 is chosen according to the second distribution also occurs with negligible probability when X_0 is chosen according to the first distribution. For these reasons, we change our security goal from an "indistinguishability-based" one (namely, requiring that, given the transcript, the real session key is indistinguishable from uniform) to an "unpredictability-based" one (namely, given the transcript, it should be infeasible to compute the real session key). In the end, though, once the parties agree on an unpredictable value k they can hash it to obtain the final session key $\mathsf{sk} = H(k)$; this final value sk will be indistinguishable from uniform if H is modeled as a random oracle.

2 Preliminaries

2.1 Notation

Let $\mathbb{Z}$ be the ring of integers, and let $[N] = \{0, 1, \ldots, N-1\}$. If χ is a probability distribution over some set S, then $x_0, x_1, \ldots, x_{\ell-1} \leftarrow \chi$ denotes independently sampling each x_i from distribution χ. We let $\mathrm{Supp}(\chi) = \{x : \chi(x) \neq 0\}$. Given an event E, we use $\overline{E}$ to denote its complement. Let $\chi(E)$ denote the probability that event E occurs under distribution χ. Given a polynomial p_i, let $(p_i)_j$ denote the jth coefficient of p_i. Let $\log(X)$ denote $\log_2(X)$, and $\exp(X)$ denote e^X.

2.2 Ring Learning with Errors

Informally, the (decisional) version of the Ring Learning with Errors (Ring-LWE) problem is: for some secret ring element s, distinguish many random "noisy ring products" with s from elements drawn uniform from the ring. More precisely, the Ring-LWE problem is parameterized by (R, q, χ, ℓ) as follows:

1. R is a ring, typically written as a polynomial quotient ring $R = \mathbb{Z}[X]/(f(X))$ for some irreducible polynomial $f(X)$ in the indeterminate X. In this paper, we restrict to the case of that $f(X) = X^n + 1$ where n is a power of 2. In later sections, we let R be parameterized by n.
2. q is a modulus defining the quotient ring $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$. We restrict to the case that q is prime and $q = 1 \mod 2n$.
3. $\chi = (\chi_s, \chi_e)$ is a pair of noise distributions over R_q (with χ_s the *secret key* distribution and χ_e the *error* distribution) that are concentrated on "short" elements, for an appropriate definition of "short" (e.g., the Euclidean distance metric on the integer-coefficients of the polynomials s or e drawn from R_q); and
4. ℓ is the number of samples provided to the adversary.

Formally, the Ring-LWE problem is to distinguish between ℓ samples independently drawn from one of two distributions. The first distribution is generated by fixing a random secret $s \leftarrow \chi_s$ then outputting

$$(a_i, b_i = s \cdot a_i + e_i) \in R_q \times R_q,$$

for $i \in [\ell]$, where each $a_i \in R_q$ is drawn uniformly at random and each $e_i \leftarrow \chi_e$ is drawn from the error distribution. For the second distribution, each sample $(a_i, b_i) \in R_q \times R_q$ is simply drawn uniformly at random.

Let A_{n,q,χ_s,χ_e} be the distribution that outputs the Ring-LWE sample $(a_i, b_i = s \cdot a_i + e_i)$ as above. We denote by $\mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_s,\chi_e,\ell}(\mathcal{B})$ the advantage of algorithm $\mathcal{B}$ in distinguishing distributions $A^{\ell}_{n,q,\chi_s,\chi_e}$ and $\mathcal{U}^{\ell}(R_q^2)$.

We define $\mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_s,\chi_e,\ell}(t)$ to be the maximum advantage of any adversary running in time t. Note that in later sections, we write as $\mathsf{Adv}_{n,q,\chi,\ell}$ when $\chi = \chi_s = \chi_e$ for simplicity.

The Ring-LWE Noise Distribution. The noise distribution χ (here we assume $\chi_s = \chi_e$, though this is not necessary) is usually a discrete Gaussian distribution on $R_q^{\vee}$ or in our case R_q (see [18] for details of the distinction, especially for concrete implementation purposes). Formally, in case of power of two cyclotomic rings, the discrete Gaussian distribution can be sampled by drawing each coefficient independently from the 1-dimensional discrete Gaussian distribution over $\mathbb{Z}$ with parameter σ, which is supported on $\{x \in \mathbb{Z} : -q/2 \leq x \leq q/2\}$ and has density function

$$D_{\mathbb{Z}_q,\sigma}(x) = \frac{e^{\frac{-\pi x^2}{\sigma^2}}}{\sum_{x=-\infty}^{\infty} e^{\frac{-\pi x^2}{\sigma^2}}}.$$

2.3 Rényi divergence

The Rényi divergence (RD) is a measure of closeness of two probability distributions. For any two discrete probability distributions P and Q such that $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$, we define the Rényi divergence of order 2 as

$$\mathrm{RD}_2(P||Q) = \sum_{x \in \mathrm{Supp}(P)} \frac{P(x)^2}{Q(x)}.$$

Rényi divergence has a probability preservation property that can be considered the multiplicative analogues of statistical distance.

Proposition 1. *Given discrete distributions P and Q with $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$, let $E \in \mathrm{Supp}(Q)$ be an arbitrary event. We have*

$$Q(E) \geq P(E)^2 / \mathrm{RD}_2(P||Q).$$

This property implies that as long as $\mathrm{RD}_2(P||Q)$ is bounded by $\mathrm{poly}(\lambda)$, any event E that occurs with negligible probability $Q(E)$ under distribution Q also occurs with negligible probability $P(E)$ under distribution P. We refer to [27, 26] for the formal proof.

Theorem 2.1 ([7]). *Fix $m, q \in \mathbb{Z}$, a bound B, and the 1-dimensional discrete Gaussian distribution $D_{\mathbb{Z}_q,\sigma}$ with parameter σ such that $B < \sigma < q$. Moreover, let $e \in \mathbb{Z}$ be such that $|e| \leq B$. If $\sigma = \Omega(B\sqrt{m/\log \lambda})$, then*

$$\mathrm{RD}_2((e + D_{\mathbb{Z}_q,\sigma})^m || D^m_{\mathbb{Z}_q,\sigma}) \leq \exp(2\pi m (B/\sigma)^2) = \mathrm{poly}(\lambda),$$

where X^m denotes m independent samples from X.

2.4 Generic Key Reconciliation Mechanism

In this subsection, we define a generic, one round, two-party key reconciliation mechanism which allows both parties to derive the same key from an approximately agreed upon ring element. A key reconciliation mechanism KeyRec consists of two algorithms recMsg and recKey, parameterized by security parameter 1^λ as well as β_{Rec}. In this context, Alice and Bob hold "close" keys – b_A and b_B, respectively – and wish to generate a shared key k so that $k = k_A = k_B$. The abstract mechanism KeyRec is defined as follows:

1. Bob computes $(K, k_B) = \mathsf{recMsg}(b_B)$ and sends the reconciliation message K to Alice.
2. Once receiving K, Alice computes $k_A = \mathsf{recKey}(b_A, K) \in \{0,1\}^\lambda$.

Correctness. Given $b_A, b_B \in R_q$, if each coefficient of $b_B - b_A$ is bounded by β_{Rec} – namely, $|b_B - b_A| \leq \beta_{\mathsf{Rec}}$ – then it is guaranteed that $k_A = k_B$.

Security. A key reconciliation mechanism KeyRec is secure if the subsequent two distribution ensembles are computationally indistinguishable. (First, we describe a simple, helper distribution.)

$\mathsf{Exe}_{\mathsf{KeyRec}}(\lambda)$: A draw from this helper distribution is performed by initiating the key reconciliation protocol among two honest parties and outputting (K, k_B); i.e. the reconciliation message K and (Bob's) key k_B of the protocol execution.

We denote by $\mathsf{Adv}_{\mathsf{KeyRec}}(\mathcal{B})$ the advantage of adversary $\mathcal{B}$ distinguishing the distributions below.

$$\{(K, k_B) : b_B \leftarrow \mathcal{U}(R_q), (K, k_B) \leftarrow \mathsf{Exe}_{\mathsf{KeyRec}}(\lambda, b_B)\}_{\lambda\in\mathbb{N}},$$
$$\{(K, k') : b_B \leftarrow \mathcal{U}(R_q), (K, k_B) \leftarrow \mathsf{Exe}_{\mathsf{KeyRec}}(\lambda, b_B), k' \leftarrow U_\lambda\}_{\lambda\in\mathbb{N}},$$

where U_λ denotes the uniform distribution over λ bits.

We define $\mathsf{Adv}_{\mathsf{KeyRec}}(t)$ to be the maximum advantage of any adversary running in time t.

Key reconciliation mechanisms from the literature. The notion of key reconciliation was first introduced by Ding et al. [19]. in his work on two-party, lattice-based key exchange. It was later used in several important works on two-party key exchange, including [28, 32, 4].

In the key reconciliation mechanisms of Peikert [28], Zhang et al. [32] and Alkim et al. [4], the initiating party sends a small amount of information about its secret, b_B, to the other party. This information is enough to allow the two parties to agree upon the same key $k = k_A = k_B$, while revealing no information about k to an eavesdropper. When instantiating our GKE protocol with this type of key reconciliation (specifically, one of [28, 32, 4]), our final GKE protocol is "contributory," in the sense that all parties contribute entropy towards determining the final key.

Another method for the two parties to agree upon the same joint key $k = k_A = k_B$, given that they start with keys b_A, b_B that are "close," was first introduced in [3] (we refer to their technique as a key reconciliation mechanism, although it is technically not referred to as such in the literature). Here, the initiating party uses its private input to generate a Regev-style encryption of a random bit string k_B of its choice under secret key b_B. and then sends to the other party, who decrypts with its approximate secret key b_A to obtain k_A. Due to the inherent robustness to noise of Regev-style encryption, it is guaranteed that $k = k_A = k_B$ with all but negligible probability. Instantiating our GKE protocol with this type of key reconciliation (specifically, that in [3]) is also possible, but does not lead to the preferred "contributory GKE," since the initiating party's entropy completely determines the final group key.

3 Group Key Exchange Security Model

A group key-exchange protocol allows a session key to be established among $N > 2$ parties. Following prior work [23, 14, 12, 13], we will use the term group key exchange (GKE) to denote a protocol secure against a *passive* (eavesdropping) adversary and will use the term *group authenticated key exchange* (GAKE) to denote a protocol secure against an *active* adversary, who controls all communication channels. Fortunately, the work of Katz and Yung [23] presents a compiler that takes any GKE protocol and transforms it into a GAKE protocol. The underlying tool required for this transform is any digital signature scheme which is strongly unforgeable under adaptive chosen message attack (EUF-CMA). We may thus focus our attention on achieving GKE in the remainder of this work.

In GKE, the adversary gets to see a single transcript generated by an execution of the GKE protocol. Given the transcript, the adversary must distinguish the real key from a fake key that is generated uniformly at random and independently of the transcript.

Formally, for security parameter $\lambda \in \mathbb{N}$, we define the following distribution:

$\mathsf{Execute}_{\Pi}^{\mathcal{O}_H}(\lambda)$: A draw from this distribution is performed by sampling a classical random oracle $\mathcal{H}$ from distribution $\mathcal{O}_H$, initiating the GKE protocol Π among N honest parties with security parameter λ relative to $\mathcal{H}$, and outputting $(\mathsf{trans}, \mathsf{sk})$—the transcript trans and key sk of the protocol execution.

Consider the following distributions:

$$\{(\mathsf{trans}, \mathsf{sk}) : (\mathsf{trans}, \mathsf{sk}) \leftarrow \mathsf{Execute}_{\Pi}^{\mathcal{O}_H}(\lambda)\}_{\lambda \in \mathbb{N}},$$
$$\{(\mathsf{trans}, \mathsf{sk}') : (\mathsf{trans}, \mathsf{sk}) \leftarrow \mathsf{Execute}_{\Pi}^{\mathcal{O}_H}(\lambda), \mathsf{sk}' \leftarrow U_\lambda\}_{\lambda \in \mathbb{N}},$$

where U_λ denotes the uniform distribution over λ bits. Let $\mathsf{Adv}^{\mathsf{GKE}, \mathcal{O}_H}(\mathcal{A})$ denote the advantage of adversary $\mathcal{A}$, with classical access to the sampled oracle $\mathcal{H}$, distinguishing the distributions above.

To enable a concrete security analysis, we define $\mathsf{Adv}^{\mathsf{GKE}, \mathcal{O}_H}(t, q_{\mathcal{O}_H})$ to be the maximum advantage of any adversary running in time t and making at most $q_{\mathcal{O}_H}$

queries to the random oracle. Security holds even if the adversary sees multiple executions by a hybrid argument.

In the next section we will define our GKE scheme and prove that it satisfies the notion of GKE.

4 A Group Key-Exchange Protocol

In this section, we present our group key exchange construction, GKE, which runs key reconciliation protocol KeyRec as a subroutine. Let KeyRec be parametrized by β_{Rec}. The protocol has two security parameters λ and ρ. λ is the computational security parameter, which is used in the security proof. ρ is the statistical security parameter, which is used in the correctness proof. σ_1, σ_2 are parameters of discrete Gaussian distributions. In this setting, N players $P_0, \ldots, P_{N-1}$ plan to generate a shared session key. The players' indices are taken modulo N.

The structure of the protocol is as follows: All parties agree on "close" keys $b_0 \approx \cdots \approx b_{N-1}$ after the second round. Player $N-1$ then initiates a key reconciliation protocol to allow all users to agree on the same key $k = k_0 = \cdots = k_{N-1}$. Since we are only able to prove that k is difficult to compute for an eavesdropping adversary (but may not be indistinguishable from random), we hash k using random oracle $\mathcal{H}$ to get the final shared key sk.

Public parameter: $R_q = \mathbb{Z}_q[x]/(x^n+1)$, $a \leftarrow \mathcal{U}(R_q)$.

Round 1: Each player P_i samples $s_i, e_i \leftarrow \chi_{\sigma_1}$ and broadcasts $z_i = as_i + e_i$.
Round 2: Player P_0 samples $e'_0 \leftarrow \chi_{\sigma_2}$ and each of the other players P_i samples $e'_i \leftarrow \chi_{\sigma_1}$, broadcasts $X_i = (z_{i+1} - z_{i-1})s_i + e'_i$.
Round 3: Player P_{N-1} proceeds as follows:

1. Samples $e''_{N-1} \leftarrow \chi_{\sigma_1}$ and computes $b_{N-1} = z_{N-2}Ns_{N-1} + e''_{N-1} + X_{N-1} \cdot (N-1) + X_0 \cdot (N-2) + \cdots + X_{N-3}$.
2. Computes $(K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1})$ and broadcasts K_{N-1}.
3. Obtains session key $\mathsf{sk}_{N-1} = \mathcal{H}(k_{N-1})$.

Key Computation: Each player P_i (except P_{N-1}) proceeds as follows:

1. Computes $b_i = z_{i-1}Ns_i + X_i \cdot (N-1) + X_{i+1} \cdot (N-2) + \cdots + X_{i+N-2}$.
2. Computes $k_i = \mathsf{recKey}(b_i, K_{N-1})$, and obtains session key $\mathsf{sk}_i = \mathcal{H}(k_i)$.

4.1 Correctness

The following claim states that each party derives the same session key sk_i, with all but negligible probability, as long as $\chi_{\sigma_1}, \chi_{\sigma_2}$ satisfy the constraint $(N^2+2N) \cdot \sqrt{n}\rho^{3/2}\sigma_1^2 + (\frac{N^2}{2}+1)\sigma_1 + (N-2)\sigma_2 \le \beta_{\mathsf{Rec}}$, where β_{Rec} is the parameter from the KeyRec protocol.

Theorem 4.1. *Given β_{Rec} as the parameter of* KeyRec *protocol, $N, n, \rho, \sigma_1, \sigma_2$ as parameters of GKE protocol Π, as long as $(N^2+2N) \cdot \sqrt{n}\rho^{3/2}\sigma_1^2 + (\frac{N^2}{2}+1)\sigma_1 + (N-2)\sigma_2 \le \beta_{\mathsf{Rec}}$ is satisfied, if all players honestly execute the group key exchange protocol described above, then each player derives the same key as input of $\mathcal{H}$ with probability $1 - 2 \cdot 2^{-\rho}$.*

Proof. We refer to Section A of Appendix for the detailed proof. □

5 Security Proof

The following theorem shows that the protocol Π is a passively secure group key-exchange protocol in random oracle model based on Ring-LWE assumption.

Theorem 5.1. *If the parameters in the group key exchange protocol Π satisfy the constraints $2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1 \le \beta_{\mathsf{Rényi}}$ and $\sigma_2 = \Omega(\beta_{\mathsf{Rényi}}\sqrt{n/\log\lambda})$, and if $\mathcal{H}$ is modeled as a classical random oracle, then for any algorithm $\mathcal{A}$ running in time t, making at most q queries to the random oracle, the maximum advantage of $\mathcal{A}$ in breaking GKE security is as follows:*

$$\mathsf{Adv}_{\Pi}^{\mathsf{GKE},\mathcal{O}_H}(t,\mathsf{q}) \le 2^{-\lambda+1} + \sqrt{\left(N \cdot \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1) + \mathsf{Adv}_{\mathsf{KeyRec}}(t_2) + \frac{\mathsf{q}}{2^\lambda}\right) \cdot \frac{\exp\left(2\pi n \left(\beta_{\mathsf{Rényi}}/\sigma_2\right)^2\right)}{1-2^{-\lambda+1}}},$$

where $t_1 = t + \mathcal{O}(N)\cdot t_{\mathsf{ring}}, t_2 = t + \mathcal{O}(N)\cdot t_{\mathsf{ring}}$ and where t_{ring} is defined as the (maximum) time required to perform operations in R_q.

Proof. Consider the joint distribution of (T,sk), where $\mathsf{T} = (\{z_i\},\{X_i\},K_{k-1})$ is the transcript of an execution of the protocol Π, and k is the final shared session key. The distribution of (T,sk) is denoted as Real. Proceeding via a sequence of experiments, we will show that under the Ring-LWE assumption, if an efficient adversary queries the random oracle on input k_{N-1} in the Ideal experiment (to be formally defined) with at most negligible probability, then it also queries the random oracle on input k_{N-1} in the Real experiment with at most negligible probability.

Furthermore, in Ideal, the input k_{N-1} to the random oracle is uniform random, which means that the adversary has $\mathsf{negl}(\lambda)$ probability of guessing k_{N-1} in Ideal when $\mathsf{q} = \mathrm{poly}(\lambda)$. Finally, we argue that the above is sufficient to prove the GKE security of the scheme, because in the random oracle model, the output of the random oracle on k_{N-1} – i.e. the agreed upon key – looks uniformly random to an adversary who does not query k_{N-1}. We now proceed with the formal proof.

Let Query be the event that k_{N-1} is among the adversary $\mathcal{A}$'s random oracle queries and denote by $\Pr_i[\mathsf{Query}]$ the probability that event Query happens in *Experiment i*. Note that we let $e'_0 = \hat{e}_0$ in order to distinguish this from the other e'_i's sampled from a different distribution.

Experiment 0. This is the original experiment. In this experiment, the distribution of (T,sk) is as follows, denoted Real :

$$
\mathsf{Real} := \left\{ \begin{array}{l} a \leftarrow R_q; s_0, s_1, \ldots, s_{N-1}, e_0, e_1, \ldots, e_{N-1} \leftarrow \chi; \\ z_0 = as_0 + e_0, z_1 = as_1 + e_1, \ldots, z_{N-1} = as_{N-1} + e_{N-1}; \\ e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; \hat{e}_0 \leftarrow \chi_{\sigma_2}; \\ X_0 = (z_1 - z_{N-1})s_0 + \hat{e}_0, X_1 = (z_2 - z_0)s_1 + e'_1, \ldots, \\ X_{N-1} = (z_0 - z_{N-2})s_{N-1} + e'_{N-1}; e''_{N-1} \leftarrow \chi_{\sigma_1}; \\ b_{N-1} = z_{N-2}Ns_{N-1} + e''_{N-1} + X_{N-1} \cdot (N-1) + \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\ (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}). \end{array} : (\mathsf{T}, \mathsf{sk}) \right\}
$$

Since $\Pr[\mathcal{A} \text{ succeeds}] = \frac{1}{2} + \mathsf{Adv}_{\Pi}^{\mathsf{GKE},\mathcal{O}_H}(t, \mathsf{q}) = \Pr_0[\mathsf{Query}] \cdot 1 + \Pr_0(\overline{\mathsf{Query}}) \cdot \frac{1}{2}$, we have

$$
\mathsf{Adv}_{\Pi}^{\mathsf{GKE},\mathcal{O}_H}(t, \mathsf{q}) \leq \Pr_0[\mathsf{Query}]. \tag{1}
$$

In the remainder of the proof, we focus on bounding $\Pr_0[\mathsf{Query}]$.

Experiment 1. In this experiment, X_0 is replaced by $X'_0 = -\sum_{i=1}^{N-1} X_i + \hat{e}_0$. The remainder of the experiment is exactly the same as *Experiment 0*. The corresponding distribution of $(\mathsf{T}, \mathsf{sk})$ is as follows, denoted Dist_1:

$$
\mathsf{Dist}_1 := \left\{ \begin{array}{l} a \leftarrow \mathcal{U}(R_q); s_0, s_1, \ldots, s_{N-1}, e_0, e_1, \ldots, e_{N-1} \leftarrow \chi_{\sigma_1}; \\ z_0 = as_0 + e_0, z_1 = as_1 + e_1, \ldots, z_{N-1} = as_{N-1} + e_{N-1}; \\ e'_0, e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; \hat{e}_0 \leftarrow \chi_{\sigma_2} \\ X'_0 = -\displaystyle\sum_{i=1}^{N-1} X_i + \hat{e}_0, X_1 = (z_2 - z_0)s_1 + e'_1, \ldots, \\ X_{N-1} = (z_0 - z_{N-2})s_{N-1} + e'_{N-1}; e''_{N-1} \leftarrow \chi_{\sigma_1}; \\ b_{N-1} = z_{N-2}Ns_{N-1} + e''_{N-1} + X_{N-1} \cdot (N-1) + \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\ (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}). \end{array} : (\mathsf{T}, \mathsf{sk}) \right\}
$$

Claim. Given $a \leftarrow \mathcal{U}(R_q)$, $s_0, s_1, \ldots, s_{N-1}, e_0, e_1, \ldots, e_{N-1}, e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}$, $\hat{e}_0 \leftarrow \chi_{\sigma_2}$, $X_0 = (z_1 - z_{N-1})s_0 + \hat{e}_0$, $X'_0 = -\sum_{i=1}^{N-1} X_i + \hat{e}_0$, where $R_q, \chi_{\sigma_1}, \chi_{\sigma_2}$, $z_1, z_{N-1}, X_1, \ldots, X_{N-1}$ are defined as above, and the constraint $2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1 \leq \beta_{\mathsf{Rényi}}$ is satisfied, we have

$$
\Pr_0[\mathsf{Query}] \leq \sqrt{\Pr_1[\mathsf{Query}] \cdot \frac{\exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2)}{1 - 2^{-\lambda+1}}} + 2^{-\lambda+1}. \tag{2}
$$

Proof. Let $\mathsf{Error} = \sum_{i=0}^{N-1}(s_i e_{i+1} + s_i e_{i-1}) + \sum_{i=1}^{N-1} e'_i$. We begin by showing that the absolute value of each coefficient of Error is bounded by $\beta_{\mathsf{Rényi}}$ with all but negligible probability. Then by adding a "bigger" error $\hat{e}_0 \leftarrow \chi_{\sigma_2}$, the small difference between distributions $\mathsf{Error} + \chi_{\sigma_2}$ (corresponding to Experiment 0) and χ_{σ_2} (corresponding to Experiment 1) can be "washed" away by applying Theorem 2.1.

For all coefficient indices j, note that $|\mathsf{Error}_j| = |(\sum_{i=0}^{N-1}(s_i e_{i+1} + s_i e_{i-1}) + \sum_{i=1}^{N-1} e'_i)_j|$. Let bound_λ denote the event that for all i and all coordinate indices j, $|(s_i)_j| \leq c\sigma_1$, $|(e_i)_j| \leq c\sigma_1$, $|(e'_i)_j| \leq c\sigma_1$,$|(e''_{N-1})_j| \leq c\sigma_1$, and $|(\hat{e}_0)_j| \leq c\sigma_2$, where $c = \sqrt{\frac{2\lambda}{\log e}}$. By replacing ρ with λ in Lemma A.1 and Lemma A.2 and by a union bound, we have – conditioned on bound_λ – that $|\mathsf{Error}_j| \leq 2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1$ for all j, with probability at least $1 - 2N \cdot 2n2^{-2\lambda}$. Since, under the assumption that $4Nn \leq 2^\lambda$, we have that $\Pr[\mathsf{bound}_\lambda] \geq 1 - 2^{-\lambda}$, we conclude that

$$\Pr[|\mathsf{Error}_j| \leq \beta_{\mathsf{Rényi}}, \forall j] \geq 1 - 2^{-\lambda+1}. \tag{3}$$

For a fixed $\mathsf{Error} \in R_q$, we denote by D_1 the distribution of $\mathsf{Error} + \chi_{\sigma_2}$ and note that D_1, χ_{σ_2} are n-dimension distributions.

Since $\sigma_2 = \Omega(\beta_{\mathsf{Rényi}}\sqrt{n/\log\lambda})$, assuming that for all j, $|\mathsf{Error}_j| \leq \beta_{\mathsf{Rényi}}$, by Theorem 2.1, we have

$$\mathrm{RD}_2(D_1 || \chi_{\sigma_2}) \leq \exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2) = \mathrm{poly}(\lambda). \tag{4}$$

Then it is straightforward to verify that the distribution of X_0 in *Experiment 0* is

$$as_1 s_0 - as_{N-1}s_0 - \sum_{i=0}^{N-1}(e_{i+1}s_i + e_{i-1}s_i) - \sum_{i=1}^{N-1} e'_i + D_1,$$

and the distribution of X'_0 in *Experiment 1* is

$$as_1 s_0 - as_{N-1}s_0 - \sum_{i=0}^{N-1}(e_{i+1}s_i + e_{i-1}s_i) - \sum_{i=1}^{N-1} e'_i + \chi_{\sigma_2}.$$

In addition, the remaining part of Dist_1 is identical to Real. Therefore we may view Real in *Experiment 0* as a function of a random variable sampled from D_1 and take Dist_1 in *Experiment 1* as a function of a random variable sampled from χ_{σ_2}.

Recall that Query is the event that k_{N-1} is contained in the set of random oracle queries issued by adversary $\mathcal{A}$. We denote by Xbound the event that $|\mathsf{Error}_j| \leq \beta_{\mathsf{Rényi}}, \forall j$. Note that computation of Error_j is available in both *Experiment 0* and *Experiment 1*. We denote by $\Pr_0[\mathsf{Xbound}]$ (resp. $\Pr_1[\mathsf{Xbound}]$) the probability that event Xbound occurs in *Experiment 0* (resp. *Experiment 1*) and define $\Pr_0[\overline{\mathsf{Xbound}}], \Pr_1[\overline{\mathsf{Xbound}}]$ analogously. Let Real' (resp. Dist'_1) denote the random variable Real (resp. Dist_1), conditioned on the event Xbound. Therefore, we have

$$
\begin{aligned}
\Pr_0[\mathsf{Query}] &= \Pr_0[\mathsf{Query}|\mathsf{Xbound}] \cdot \Pr_0[\mathsf{Xbound}] + \Pr_0[\mathsf{Query}|\overline{\mathsf{Xbound}}] \cdot \Pr_0[\overline{\mathsf{Xbound}}] \\
&\leq \Pr_0[\mathsf{Query}|\mathsf{Xbound}] + \Pr_0[\overline{\mathsf{Xbound}}] \\
&\leq \Pr_0[\mathsf{Query}|\mathsf{Xbound}] + 2^{-\lambda+1} \\
&\leq \sqrt{\Pr_1[\mathsf{Query}|\mathsf{Xbound}] \cdot \mathrm{RD}_2(\mathsf{Real}'||\mathsf{Dist}_1')} + 2^{-\lambda+1} \\
&\leq \sqrt{\Pr_1[\mathsf{Query}|\mathsf{Xbound}] \cdot \mathrm{RD}_2(D_1||\chi_{\sigma_2})} + 2^{-\lambda+1} \\
&\leq \sqrt{\Pr_1[\mathsf{Query}|\mathsf{Xbound}] \cdot \exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2)} + 2^{-\lambda+1} \\
&\leq \sqrt{\left(\Pr_1[\mathsf{Query} \cdot \frac{\exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2)}{\Pr_1[\mathsf{Xbound}]}\right.} + 2^{-\lambda+1} \\
&\leq \sqrt{\left(\Pr_1[\mathsf{Query}] \cdot \frac{\exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2)}{1-2^{-\lambda+1}}\right.} + 2^{-\lambda+1},
\end{aligned}
$$

where the second and last inequalities follow from (3), the third inequality follows from Proposition 1 and the fifth inequality follows from (4). □

In Section B of the Appendix, we show that

$$\Pr_1[\mathsf{Query}] \leq \left(N \cdot \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1) + \mathsf{Adv}_{\mathsf{KeyRec}}(t_2) + \frac{\mathsf{q}}{2^\lambda}\right)\Bigg($$

which concludes the proof of Theorem 5.1. □

5.1 Parameter Constraints

Beyond the parameter settings recommended for instantiating Ring-LWE with security parameter λ, parameters $N, n, \sigma_1, \sigma_2, \lambda, \rho$ of the protocol above are also required to satisfy the following inequalities:

$$(N^2 + 2N) \cdot \sqrt{n}\rho^{3/2}\sigma_1^2 + (\frac{N^2}{2} + 1)\sigma_1 + (N-2)\sigma_2 \leq \beta_{\mathsf{Rec}} \quad \text{(Correctness)} \tag{5}$$

$$2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1 \leq \beta_{\mathsf{Rényi}} \quad \text{(Security)} \tag{6}$$

$$\sigma_2 = \Omega(\beta_{\mathsf{Rényi}}\sqrt{n/\log\lambda}) \quad \text{(Security)} \tag{7}$$

We comment that once the ring, the noise distributions, and the security parameters λ, ρ are fixed, the maximum number of parties is fixed.

6 Acknowledgments

This material is based on work performed under financial assistance award 70NANB15H328 from the U.S. Department of Commerce, National Institute of Standards and Technology. Work by Dana Dachman-Soled was additionally supported in part by NSF grants #CNS-1840893 and #CNS-1453045, and by a research partnership award from Cisco.

References

1. Michel Abdalla, Emmanuel Bresson, Olivier Chevassut, and David Pointcheval. Password-based group key exchange in a constant number of rounds. In *9th Intl. Conference on Theory and Practice of Public Key Cryptography (PKC)*, volume 3958 of *Lecture Notes in Computer Science*, pages 427–442. Springer, 2006.
2. Michel Abdalla and David Pointcheval. A scalable password-based group key exchange protocol in the standard model. In *Advances in Cryptology—Asiacrypt 2006*, volume 4284 of *Lecture Notes in Computer Science*, pages 332–347. Springer, 2006.
3. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. NewHope without reconciliation. Cryptology ePrint Archive, Report 2016/1157, 2016. http://eprint.iacr.org/2016/1157.
4. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange—a new hope. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 327–343, Austin, TX, 2016. USENIX Association.
5. Klaus Becker and Uta Wille. Communication complexity of group key distribution. In *Proceedings of the 5th ACM Conference on Computer and Communications Security*, CCS '98, pages 1–6, New York, NY, USA, 1998.
6. Mihir Bellare and Phillip Rogaway. Provably secure session key distribution: The three party case. In *27th Annual ACM Symposium on Theory of Computing*, pages 57–66, Las Vegas, NV, USA, May 29 – June 1, 1995. ACM Press.
7. Andrej Bogdanov, Siyao Guo, Daniel Masny, Silas Richelson, and Alon Rosen. On the hardness of learning with rounding over small modulus. In *Theory of Cryptography Conference*, pages 209–224. Springer, 2016.
8. Jens-Matthias Bohli, Maria Isabel Gonzalez Vasco, and Rainer Steinwandt. Password-authenticated constant-round group key establishment with a common reference string. Cryptology ePrint Archive, Report 2006/214, 2006. http://eprint.iacr.org/2006/214.
9. Jens-Matthias Bohli, María Isabel González Vasco, and Rainer Steinwandt. Secure group key establishment revisited. *International Journal of Information Security*, 6(4):243–254, Jul 2007.
10. Dan Boneh, Darren Glass, Daniel Krashen, Kristin Lauter, Shahed Sharif, Alice Silverberg, Mehdi Tibouchi, and Mark Zhandry. Multiparty non-interactive key exchange and more from isogenies on elliptic curves. *arXiv preprint arXiv:1807.03038*, 2018.
11. Emmanuel Bresson and Dario Catalano. Constant round authenticated group key agreement via distributed computation. In Feng Bao, Robert Deng, and Jianying Zhou, editors, *PKC 2004: 7th Intl. Workshop on Theory and Practice in Public Key Cryptography*, volume 2947 of *Lecture Notes in Computer Science*, pages 115–129, Singapore, March 1–4, 2004. Springer.
12. Emmanuel Bresson, Olivier Chevassut, and David Pointcheval. Provably authenticated group Diffie-Hellman key exchange – the dynamic case. In Colin Boyd, editor, *Advances in Cryptology—Asiacrypt 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 290–309, Gold Coast, Australia, December 9–13, 2001. Springer.
13. Emmanuel Bresson, Olivier Chevassut, and David Pointcheval. Dynamic group Diffie-Hellman key exchange under standard assumptions. In Lars R. Knudsen, editor, *Advances in Cryptology—Eurocrypt 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 321–336, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer.

14. Emmanuel Bresson, Olivier Chevassut, David Pointcheval, and Jean-Jacques Quisquater. Provably authenticated group Diffie-Hellman key exchange. In *ACM CCS 01: 8th Conference on Computer and Communications Security*, pages 255–264, Philadelphia, PA, USA, November 5–8, 2001. ACM Press.
15. Mike Burmester and Yvo Desmedt. A secure and efficient conference key distribution system (extended abstract). In Alfredo De Santis, editor, *Advances in Cryptology—Eurocrypt'94*, volume 950 of *Lecture Notes in Computer Science*, pages 275–286. Springer, 1995.
16. Mike Burmester and Yvo Desmedt. A secure and scalable group key exchange system. *Information Processing Letters*, 94(3):137–143, May 2005.
17. Kyu Young Choi, Jung Yeon Hwang, and Dong Hoon Lee. Efficient ID-based group key agreement with bilinear maps. In Feng Bao, Robert Deng, and Jianying Zhou, editors, *PKC 2004: 7th Intl. Workshop on Theory and Practice in Public Key Cryptography*, volume 2947 of *Lecture Notes in Computer Science*, pages 130–144, Singapore, March 1–4, 2004. Springer.
18. Eric Crockett and Chris Peikert. Challenges for ring-LWE. Cryptology ePrint Archive, Report 2016/782, 2016. http://eprint.iacr.org/2016/782.
19. Jintai Ding, Xiang Xie, and Xiaodong Lin. A simple provably secure key exchange scheme based on the learning with errors problem. Cryptology ePrint Archive, Report 2012/688, 2012. http://eprint.iacr.org/2012/688.
20. Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
21. I. Ingemarsson, D. Tang, and C. Wong. A conference key distribution system. *IEEE Trans. Inf. Theor.*, 28(5):714–720, September 1982.
22. Jonathan Katz and Ji Sun Shin. Modeling insider attacks on group key-exchange protocols. In *Proceedings of the 12th ACM Conference on Computer and Communications Security*, CCS '05, pages 180–189, New York, NY, USA, 2005. ACM.
23. Jonathan Katz and Moti Yung. Scalable protocols for authenticated group key exchange. In Dan Boneh, editor, *Advances in Cryptology—Crypto 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 110–125, Santa Barbara, CA, USA, 2003. Springer.
24. Jonathan Katz and Moti Yung. Scalable protocols for authenticated group key exchange. *Journal of Cryptology*, 20(1):85–113, 2007.
25. Yongdae Kim, Adrian Perrig, and Gene Tsudik. Simple and fault-tolerant key agreement for dynamic collaborative groups. In *Proceedings of the 7th ACM Conference on Computer and Communications Security*, CCS '00, pages 235–244, New York, NY, USA, 2000.
26. Adeline Langlois, Damien Stehlé, and Ron Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology—Eurocrypt 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 239–256, Copenhagen, Denmark, May 11–15, 2014. Springer.
27. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology—Eurocrypt 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23, French Riviera, May 30 – June 3, 2010. Springer.
28. Chris Peikert. Lattice cryptography for the internet. Cryptology ePrint Archive, Report 2014/070, 2014. http://eprint.iacr.org/2014/070.
29. D. G. Steer and L. Strawczynski. A secure audio teleconference system. In *MILCOM 88, 21st Century Military Communications - What's Possible?'. Conference record. Military Communications Conference*, Oct 1988.

30. M. Steiner, G. Tsudik, and M. Waidner. Key agreement in dynamic peer groups. *IEEE Transactions on Parallel and Distributed Systems*, 11(8):769–780, Aug 2000.
31. Qianhong Wu, Yi Mu, Willy Susilo, Bo Qin, and Josep Domingo-Ferrer. Asymmetric group key agreement. In Antoine Joux, editor, *Advances in Cryptology—Eurocrypt 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 153–170, Cologne, Germany, April 26–30, 2009. Springer.
32. Jiang Zhang, Zhenfeng Zhang, Jintai Ding, Michael Snook, and Özgür Dagdelen. Authenticated key exchange from ideal lattices. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology—Eurocrypt 2015, Part II*, volume 9057 of *Lecture Notes in Computer Science*, pages 719–751, Sofia, Bulgaria, April 26–30, 2015. Springer.

A Correctness of the Group Key-Exchange Protocol

Theorem 4.1. *Given β_{Rec} as parameter of* KeyRec *protocol, $N, n, \rho, \sigma_1, \sigma_2$ as parameters of GKE protocol Π, $(N^2+2N)\cdot\sqrt{n}\rho^{3/2}\sigma_1^2+(\frac{N^2}{2}+1)\sigma_1+(N-2)\sigma_2 \leq \beta_{\mathsf{Rec}}$ is satisfied, if all players honestly execute the group key exchange protocol as described above, then each player derive the same key as input of $\mathcal{H}$ with probability $1-2\cdot 2^{-\rho}$.*

Proof. Given $s_i, e_i, e'_i, e''_{N-1} \leftarrow \chi_{\sigma_1}$, $\hat{e}_0 \leftarrow \chi_{\sigma_2}$ for all i as specified in protocol Π, we begin by introducing the following lemmas to analyze probabilities that each coordinate of $s_i, e_i, e'_i, e''_{N-1}, \hat{e}_0$ are "short" for all i, and conditioned on the first event, $s_i e_i$ are "short".

Lemma A.1. *Given $s_i, e_i, e'_i, e''_{N-1}, \hat{e}_0$ for all i as defined above, let* bound *denote the event that for all i and all coordinate indices j, $|(s_i)_j| \leq c\sigma_1$, $|(e_i)_j| \leq c\sigma_1$, $|(e'_i)_j| \leq c\sigma_1$,$|(e''_{N-1})_j| \leq c\sigma_1$, and $|(\hat{e}_0)_j| \leq c\sigma_2$, where $c = \sqrt{\frac{2\rho}{\pi \log e}}$, we have* $\Pr[\mathsf{bound}] \geq 1-2^{-\rho}$.

Proof. Using the fact that complementary error function $\operatorname{erfc}(x) = \frac{2}{\sqrt{\pi}}\int_x^\infty e^{-t^2}dt \leq e^{-x^2}$, we obtain

$$\begin{aligned}\Pr[|v| \geq c\sigma+1; v \leftarrow D_{\mathbb{Z}_q,\sigma}] &\leq 2\sum_{x=\lfloor c\sigma+1\rceil}^{\infty} D_{\mathbb{Z}_q,\sigma}(x) \leq \frac{2}{\sigma}\int_{c\sigma}^{\infty} e^{-\frac{\pi x^2}{\sigma^2}}dx \\ &= \frac{2}{\sqrt{\pi}}\int_{\frac{\sqrt{\pi}}{\sigma}(c\sigma)}^{\infty} e^{-t^2}dt \leq e^{-c^2\pi}.\end{aligned}$$

Note that there are $3nN$ number of coordinates sampled from distribution $D_{\mathbb{Z}_q,\sigma_1}$, and n number of coordinates sampled from distribution $D_{\mathbb{Z}_q,\sigma_2}$ in total. Assume $3nN+n \leq e^{c^2\pi/2}$, since all the coordinates are sampled independently, we bound $\Pr[\mathsf{bound}]$ as follow:

$$\begin{aligned}\Pr[\mathsf{bound}] &= \big(1-\Pr[|v| \geq c\sigma_1+1; v \leftarrow D_{\mathbb{Z}_q,\sigma_1}]\big)^{3nN} \\ &\quad \cdot \big(1-\Pr[|\hat{e}_0| \geq c\sigma_2+1; \hat{e}_0 \leftarrow D_{\mathbb{Z}_q,\sigma_2}]\big)^{n} \\ &\geq 1-(3nN+n)e^{-c^2\pi} \geq 1-e^{-c^2\pi/2} \geq 1-2^{-\rho}.\end{aligned}$$

The last inequality follows as $c = \sqrt{\frac{2\rho}{\pi \log e}}$. □

Lemma A.2. *Given $s_i, e_i, e'_i, e''_{N-1}, \hat{e}_0$ for all i as defined above, and* bound *as defined in Lemma A.1, let* $\mathsf{product}_{\mathsf{s_i,e_j}}$ *denote the event that, for all coefficient indices v, $|(s_i e_j)_v| \le \sqrt{n}\rho^{3/2}\sigma_1^2$. we have*

$$\Pr[\mathsf{product}_{\mathsf{s_i,e_j}} | \mathsf{bound}] \ge 1 - 2n \cdot 2^{-2\rho}.$$

Proof. For $t \in \{0, \ldots, n-1\}$, Let $(s_i)_t$ denote the t^{th} coefficient of $s_i \in R_q$, namely, $s_i = \sum_{t=0}^{n-1}(s_i)_t X^i$. $(e_j)_t$ is defined analogously. Since we have X^n+1 as modulo of R, it is easy to see that $(s_i e_j)_v = c_v X^v$, where $c_v = \sum_{u=0}^{n-1}(s_i)_u(e_j)^*_{v-u}$, and $(e_j)^*_{v-u} = (e_j)_{v-u}$ if $v - u \ge 0$, $(e_j)^*_{v-u} = -(e_j)_{v-u+n}$, otherwise. Thus, conditioned on $|(s_i)_t| \le c\sigma_1$ and $|(e_j)_t| \le c\sigma_1$ (for all i, j, t) where $c = \sqrt{\frac{2\rho}{\pi \log e}}$, by Hoeffding's Inequality [20], we derive

$$\Pr[|(s_i e_j)_v| \ge \delta] = \Pr\left[\left|\sum_{u=0}^{n-1}(s_i)_u(e_j)^*_{v-u}\right| \ge \delta\right] \le 2\exp\left(\frac{-2\delta^2}{n(2c^2\sigma_1^2)^2}\right),$$

as each product $(s_i)_u(e_j)^*_{v-u}$ in the sum is an independent random variable with mean 0 in the range $[-c^2\sigma_1^2, c^2\sigma_1^2]$. By setting $\delta = \sqrt{n}\rho^{3/2}\sigma_1^2$, we obtain

$$\Pr[|(s_u e_v)_i| \ge \sqrt{n}\rho^{3/2}\sigma_1^2] \le 2^{-2\rho+1}. \tag{8}$$

Finally, by Union Bound,

$$\Pr[\mathsf{product}_{\mathsf{s_i,e_j}} | \mathsf{bound}] = \Pr[|(s_i e_j)_v| \le \sqrt{n}\rho^{3/2}\sigma_1^2, \forall v] \ge 1 - 2n \cdot 2^{-2\rho}. \tag{9}$$

□

Now we begin analyzing the chance that not all parties agree on the same final key. The correctness of KeyRec guarantees that this group key exchange protocol has agreed session key among all parties $\forall i, k_i = k_{N-1}$, if $\forall j$, the j^{th} coefficient of $|b_{N-1} - b_i| \le \beta_{\mathsf{Rec}}$.

For better illustration, we first write $X_0, \ldots, X_{N-1}$ in form of linear system as follows. $\boldsymbol{X} = [X_0 \quad X_1 \quad X_2 \quad \cdots \quad X_{N-1}]^T$

$$= \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & \ldots & 0 & -1 \\ -1 & 1 & 0 & 0 & \ldots & 0 & 0 \\ 0 & -1 & 1 & 0 & \ldots & 0 & 0 \\ 0 & 0 & -1 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & & \\ 0 & 0 & 0 & 0 & \ldots & -1 & 1 \end{bmatrix}}_{\boldsymbol{M}} \left[\underbrace{\begin{bmatrix} as_0s_1 \\ as_1s_2 \\ as_2s_3 \\ as_3s_4 \\ \vdots \\ as_{N-2}s_{N-1} \\ as_{N-1}s_0 \end{bmatrix}}_{\boldsymbol{S}} + \underbrace{\begin{bmatrix} s_0e_1 - s_0e_{N-1} + e'_0 \\ s_1e_2 - s_1e_0 + e'_1 \\ s_2e_3 - s_2e_1 + e'_2 \\ s_3e_4 - s_3e_2 + e'_3 \\ \vdots \\ s_{N-2}e_{N-3} - s_{N-2}e_{N-3} + e'_{N-2} \\ s_{N-1}e_0 - s_{N-1}e_{N-2} + e'_{N-1} \end{bmatrix}}_{\boldsymbol{E}}\right]. \tag{10}$$

We denote the matrices above by $\boldsymbol{M}, \boldsymbol{S}, \boldsymbol{E}$ from left to right and have the linear system as $\boldsymbol{X} = \boldsymbol{MS} + \boldsymbol{E}$. By setting $\boldsymbol{B}_i = [i-1 \;\; i-2 \;\; \cdots \;\; 0 \;\; N-1 \;\; N-2 \;\; \cdots \;\; i]$ as a N-dimensional vector, we can then write b_i as $\boldsymbol{B}_i \cdot \boldsymbol{X} + N(as_i s_{i-1} + s_i e_{i-1}) = \boldsymbol{B}_i \boldsymbol{MS} + \boldsymbol{B}_i \boldsymbol{E} + N(as_i s_{i-1} + s_i e_{i-1})$, for $i \neq N-1$ and write b_{N-1} as $\boldsymbol{B}_{N-1}\boldsymbol{MS} + \boldsymbol{B}_{N-1}\boldsymbol{E} + N(as_{N-1}s_{N-2} + s_{N-1}e_{N-2}) + e''_{N-1}$. It is straightforward to see that, entries of $\boldsymbol{MS}$ and $Nas_i s_{i-1}$ are eliminated through the process of computing $b_{N-1} - b_i$. Thus we get

$$\begin{aligned} b_{N-1} - b_i &= (\boldsymbol{B}_{N-1} - \boldsymbol{B}_i)\,\boldsymbol{E} + N(s_{N-1}e_{N-2} - s_i e_{i-1}) + e''_{N-1} \\ = \;& (N-i-1) \cdot \left(\left(\sum_{\substack{j \in \mathbb{Z} \cap [0, i-1] \\ \text{and } j = N-1}} \left(s_j e_{j+1} - s_j e_{j-1} + e'_j \right) \right) \right) + e''_{N-1} \\ &+ (-i-1) \left(\sum_{j=i}^{N-2} \left(s_j e_{j+1} - s_j e_{j-1} + e'_j \right) \right) + N(s_{N-1}e_{N-2} - s_i e_{i-1}) \end{aligned}$$

Observe that for an arbitrary $i \in [N]$, there are at most $(N^2 + 2N)$ terms in form of $s_u e_v$, at most $N^2/2$ terms in form of e'_w where $e'_w \leftarrow \chi_{\sigma_1}$, at most $N-2$ terms of e'_0, where $e'_0 \leftarrow \chi_{\sigma_2}$, and one term in form of e''_{N-1} in any coordinate of the sum above. Let $\mathsf{product}_{\mathsf{ALL}}$ denote the event that for all the terms in form of $s_u e_v$ observed above, each coefficient of such term is bounded by $\sqrt{n}\rho^{3/2}\sigma_1^2$. By Union Bound and by assuming $2n(N^2 + 2N) \leq 2^\rho$, it is straightforward to see $\Pr[\overline{\mathsf{product}_{\mathsf{ALL}}} | \mathsf{bound}] \leq (N^2 + 2N) \cdot 2n2^{-2\rho} \leq 2^{-\rho}$.

Let bad be the event that not all parties agree on the same final key. Given the constraint $(N^2 + 2N) \cdot \sqrt{n}\rho^{3/2}\sigma_1^2 + (\frac{N^2}{2} + 1)\sigma_1 + (N-2)\sigma_2 \leq \beta_{\mathsf{Rec}}$ satisfied, we have

$$\begin{aligned} \Pr[\mathsf{bad}] &= \Pr[\mathsf{bad}|\mathsf{bound}] \cdot \Pr[\mathsf{bound}] + \Pr[\mathsf{bad}|\overline{\mathsf{bound}}] \cdot \Pr[\overline{\mathsf{bound}}] && (11) \\ &\leq \Pr[\overline{\mathsf{product}_{\mathsf{ALL}}}] \cdot 1 + 1 \cdot \Pr[\overline{\mathsf{bound}}] \leq 2 \cdot 2^{-\rho}, && (12) \end{aligned}$$

which completes the proof.

□

B Concluding the Proof of Theorem 5.1

Theorem 5.1 (Restated). *If the parameters in group key exchange protocol Π satisfy the constraints that $2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1 \leq \beta_{\mathsf{R\acute{e}nyi}}, \sigma_2 = \Omega(\beta_{\mathsf{R\acute{e}nyi}}\sqrt{n/\log\lambda})$, and $\mathcal{H}$ is modeled as a classical random oracle, then for any algorithm $\mathcal{A}$ running in time t, making at most q queries to the random oracle, the maximum advantage of $\mathcal{A}$ in breaking* GKE *security is as follows:*

$$\begin{aligned} &\mathsf{Adv}_{\Pi}^{\mathsf{GKE},\mathcal{O}_H}(t, \mathsf{q}) \leq 2^{-\lambda+1} \\ &+ \sqrt{\left(N \cdot \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1) + \mathsf{Adv}_{\mathsf{KeyRec}}(t_2) + \frac{\mathsf{q}}{2^\lambda} \right) \cdot \frac{\exp\left(2\pi n \left(\beta_{\mathsf{R\acute{e}nyi}}/\sigma_2\right)^2\right)}{1 - 2^{-\lambda+1}}}, \end{aligned}$$

where t_1 and t_2 equal to $t + \mathcal{O}(N) \cdot t_{\mathsf{ring}}$ and t_{ring} is the time to perform operations in R_q.

Proof. (Continued) Recall that *Experiment 0* is the real world experiment. We have that $\mathsf{Adv}_{\Pi}^{\mathsf{GKE},\mathcal{O}_H}(t, \mathsf{q}) \leq \Pr_0[\mathsf{Query}]$ (see Equation 1), where Query is the event that k_{N-1} is among the adversary $\mathcal{A}$'s random oracle queries and $\Pr_i[\mathsf{Query}]$ is the probability that event Query happens in *Experiment i*.

In *Experiment 1*, we switched from X_0 as sampled in the real world to $X_0' = -\sum_{i=1}^{N-1} X_i + \hat{e}_0$ and showed (see Equation 2) that

$$\Pr_0[\mathsf{Query}] \leq \sqrt{\left(\Pr_1[\mathsf{Query}] \cdot \frac{\exp(2\pi n(\beta_{\mathsf{R\acute{e}nyi}}/\sigma_2)^2)}{1 - 2^{-\lambda+1}}\right.} + 2^{-\lambda+1}.$$

Therefore, to prove the theorem, it remains to show that

$$\Pr_1[\mathsf{Query}] \leq \left(N \cdot \mathsf{Adv}_{n,q,\chi_{\sigma_1},3}^{\mathsf{RLWE}}(t_1) + \mathsf{Adv}_{\mathsf{KeyRec}}(t_2) + \frac{\mathsf{q}}{2^{\lambda}}\right).$$

We do so by considering a sequence of experiments as follows:

Experiment 2. This experiment proceeds exactly the same as *Experiment 1*, except that z_0 is generated uniformly at random, instead of being generated as an Ring-LWE instance. The corresponding distribution is as follows, denoted Dist_2:

$$\mathsf{Dist}_2 := \left\{ \begin{array}{ll} a \leftarrow \mathcal{U}(R_q); s_1, \ldots, s_{N-1}, e_1, \ldots, e_{N-1} \leftarrow \chi_{\sigma_1}; & \\ z_0 \leftarrow \mathcal{U}(R_q), z_1 = as_1 + e_1, \ldots, z_{N-1} = as_{N-1} + e_{N-1}; & \\ e_1', \ldots, e_{N-1}' \leftarrow \chi_{\sigma_1}; \hat{e}_0 \leftarrow \chi_{\sigma_2} & \\ X_0' = -\sum_{i=1}^{N-1} X_i + \hat{e}_0, X_1 = (z_2 - z_0)s_1 + e_2', \ldots, & \\ X_{N-1} = (z_0 - z_{N-2})s_{N-1} + e_{N-1}'; e_{N-1}'' \leftarrow \chi_{\sigma_1}; & : (\mathsf{T}, \mathsf{sk}) \\ b_{N-1} = z_{N-2}Ns_{N-1} + e_{N-1}'' + X_{N-1} \cdot (N-1) + & \\ \quad X_0 \cdot (N-2) + \cdots + X_{N-3}; & \\ (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}). & \end{array} \right\}$$

Bounding the difference of $|\Pr_2[\mathsf{Query}] - \Pr_1[\mathsf{Query}]|$:

Given algorithm $\mathcal{A}$ running in time t attacking Π, let $\mathcal{B}$ be an algorithm running in time t_1 that takes as input (a, z_0), generates $(\mathsf{T}, \mathsf{sk})$ based on distribution Dist_1' which is identical to Dist_1 except for (a, z_0) given as input, runs $\mathcal{A}$ as subroutine and outputs whatever $\mathcal{A}$ outputs. It is straightforward to see that if (a, z_0) is sampled from the Ring-LWE distribution $A_{n,q,\chi_{\sigma_1}}$, then Dist_1' is identical to Dist_1, and if (a, z_0) is sampled from $\mathcal{U}(R_q^2)$, Dist_1' is identical to Dist_2. Note that t_1 is equal to t plus a minor overhead for the simulation of the security experiment for $\mathcal{A}$.

Therefore we conclude that the difference of algorithm $\mathcal{A}$'s success probability in *Experiment 1* and *Experiment 2* is bounded by probability that $\mathcal{B}$ running in time t_1 distinguishes $A_{n,q,\chi_{\sigma_1}}$ from $\mathcal{U}(R_q)$ given one sample. Since $\mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1) \geq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},2}(t_1) \geq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},1}(t_1)$, for simplicity, we have

$$|\Pr_2[\mathsf{Query}] - \Pr_1[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1). \tag{13}$$

Recall that in the previous experiment, we switched z_0 to be uniformly distributed in R_q. In next two experiments, we switch z_1, X_1 to be elements uniformly distributed in R_q.

Experiment 3. the experiment proceeds exactly the same as *Experiment 2*, except for setting $z_0 = z_2 - r_1, X_1 = r_1 s_1 + e'_1$, where r_1 is sampled from $\mathcal{U}(R_q)$. The corresponding distribution is as follows, denoted as Dist_3.

Bounding the difference of $|\Pr_3[\mathsf{Query}] - \Pr_2[\mathsf{Query}]|$: Since r_1 is sampled uniformly, $z_2 - r_1$ is also a uniformly distributed random value, then we claim that *Experiment 3* is identical to *Experiment 3* up to variable substitution, namely

$$\Pr_3[\mathsf{Query}] = \Pr_2[\mathsf{Query}]. \tag{14}$$

$$\mathsf{Dist}_3 := \left\{ \begin{array}{ll} \begin{cases} a \leftarrow \mathcal{U}(R_q), r_1 \leftarrow \mathcal{U}(R_q); \\ s_1, \ldots, s_{N-1}, e_1, \ldots, e_{N-1} \leftarrow \chi_{\sigma_1}; z_0 = z_2 - r_1, \\ z_1 = as_1 + e_1, \ldots, z_{N-1} = as_{N-1} + e_{N-1}; \\ e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; \hat{e}_0 \leftarrow \chi_{\sigma_2}; X'_0 = -\sum_{i=1}^{N-1} X_i + \hat{e}_0, \\ X_1 = r_1 s_1 + e'_1, X_2 = (z_3 - z_1)s_2 + e'_2, \\ \ldots, X_{N-1} = (z_0 - z_{N-2})s_{N-1} + e'_{N-1}; \\ e''_{N-1} \leftarrow \chi_{\sigma_1}; \\ b_{N-1} = z_{N-2} N s_{N-1} + e''_{N-1} + X_{N-1} \cdot (N-1) + \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\ (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}). \end{cases} & : (\mathsf{T}, \mathsf{sk}) \end{array} \right\}$$

Experiment 4. This experiment proceeds exactly the same as *Experiment 3*, except that z_1, X_1 are uniformly distributed in R_q. The corresponding distribution is as follows, denoted as Dist_4.

$$
\mathsf{Dist}_4 := \left\{ \begin{array}{l} a, r_1 \leftarrow \mathcal{U}(R_q); s_2, \ldots, s_{N-1}, e_2, \ldots, e_{N-1} \leftarrow \chi_{\sigma_1}; \\ z_0 = z_2 - r_1, z_1 \leftarrow \mathcal{U}(R_q), z_2 = as_2 + e_2, \ldots, \\ z_{N-1} = as_{N-1} + e_{N-1}; e'_2, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; \\ \hat{e}_0 \leftarrow \chi_{\sigma_2}; X'_0 = -\sum_{i=1}^{N-1} X_i + \hat{e}_0, \\ X_1 \leftarrow \mathcal{U}(R_q), X_2 = (z_3 - z_1)s_2 + e'_2, \\ \ldots, X_{N-1} = (z_0 - z_{N-2})s_{N-1} + e'_{N-1}; e''_{N-1} \leftarrow \chi_{\sigma_1}; \\ b_{N-1} = z_{N-2}Ns_{N-1} + e''_{N-1} + X_{N-1} \cdot (N-1) + \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\ (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}). \end{array} : (\mathsf{T}, \mathsf{sk}) \right\}
$$

Bounding the difference of $|\Pr_4[\mathsf{Query}] - \Pr_3[\mathsf{Query}]|$:

Given an algorithm $\mathcal{A}$ running in time t attacking Π, let $\mathcal{B}$ be an algorithm running in time t_1 that takes as input (a, z_1), (r_1, X_1), generates $(\mathsf{T}, \mathsf{sk})$ based on distribution Dist'_3 which is identical to Dist_3 except for (a, z_1), (r_1, X_1) given as input. $\mathcal{B}$ runs $\mathcal{A}$ as a subroutine and outputs whatever $\mathcal{A}$ outputs. Note that t_1 is equal to t plus a minor overhead for the simulation of the security experiment for $\mathcal{A}$.

It is clear to see that if (a, z_1) and (r_1, X_1) are sampled from the Ring-LWE distribution $A_{n,q,\chi_{\sigma_1}}$, then Dist'_3 is identical to Dist_3. If (a, z_1) and (r_1, X_1) are sampled from $\mathcal{U}(R_q^2)$, Dist'_3 is identical to Dist_4.

Therefore we conclude that the difference of algorithm $\mathcal{A}$ successful probability in winning *Experiment 4* and *Experiment 3* is bounded by the advantage of adversary $\mathcal{B}$ running in time t_1 in distinguishing $A_{n,q,\chi_{\sigma_1}}$ from $\mathcal{U}(R_q^2)$ given two samples. Thus,

$$
|\Pr_4[\mathsf{Query}] - \Pr_3[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1). \tag{15}
$$

Experiment 5. This experiment proceeds exactly the same as *Experiment 4*, except that z_0 is sampled directly from $\mathcal{U}(R_q)$. We leave the formal definition of Dist_5 implicit for simplicity.

Bounding the difference of $|\Pr_5[\mathsf{Query}] - \Pr_4[\mathsf{Query}]|$: It is easy to see that the corresponding distribution Dist_5 is identical to Dist_4 by substituting variable z_0 for $z_2 - r_1$. Thus,

$$
\Pr_5[\mathsf{Query}] = \Pr_4[\mathsf{Query}]. \tag{16}
$$

In the case that $N \geq 3$, we present the following sequence of experiments from *Experiment 6* to *Experiment* $3N - 4$. For $i = 2, 3, \ldots, N-2$, we define three experiments *Experiment* $3i$, *Experiment* $3i + 1$, *Experiment* $3i + 2$. It is ensured that in the experiments prior to *Experiment* $3i$, we already switched

z_j, X_j for all $0 \leq j \leq i-1$. In *Experiment* $3i$, *Experiment* $3i+1$ and *Experiment* $3i+2$, we replace z_i and X_i by random elements uniformly distributed in R_q. *Experiment* $3i$, *Experiment* $3i+1$, *Experiment* $3i+2$ are formally defined as follows:

Experiment $3i$. The experiment proceeds exactly the same as *Experiment* $3i-1$, except for setting $z_{i-1} = z_{i+1} - r_i$, $X_i = r_i s_i + e_i'$, where r_1 is sampled from $\mathcal{U}(R_q)$. The corresponding distribution is as follows, denoted Dist_{3i}

$$\mathsf{Dist}_{3i} := \left\{ \begin{array}{l} a, r_i \leftarrow \mathcal{U}(R_q); s_i, \ldots, s_{N-1}, e_i, \ldots, e_{N-1} \leftarrow \chi_{\sigma_1}; \\ z_0, \ldots, z_{i-2} \leftarrow \mathcal{U}(R_q), z_{i-1} = z_{i+1} - r_i, z_i = a s_i + e_i, \\ \ldots, z_{N-1} = a s_{N-1} + e_{N-1}; e_i', \ldots, e_{N-1}' \leftarrow \chi_{\sigma_1}; \\ \hat{e}_0 \leftarrow \chi_{\sigma_2}; X_0' = -\sum_{i=1}^{N-1} X_i + \hat{e}_0, X_1, \ldots, X_{i-1} \leftarrow \mathcal{U}(R_q) \qquad : (\mathsf{T}, \mathsf{sk}) \\ X_i = r_i s_i + e_i', X_{i+1} = (z_{i+2} - z_i) s_{i+1} + e_{i+1}', \\ \ldots, X_{N-1} = (z_0 - z_{N-2}) s_{N-1} + e_{N-1}'; e_{N-1}'' \leftarrow \chi_{\sigma_1}; \\ b_{N-1} = z_{N-2} N s_{N-1} + e_{N-1}'' + X_{N-1} \cdot (N-1) + \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\ (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}). \end{array} \right\}$$

Experiment $3i+1$. This experiment proceeds exactly the same as *Experiment* $3i$, except that z_i, X_i are uniformly distributed in R_q. The corresponding distribution is as follows, denoted Dist_{3i+1}:

$$\mathsf{Dist}_{3i+1} := \left\{ \begin{array}{l} a, r_i \leftarrow \mathcal{U}(R_q); s_{i+1}, \ldots, s_{N-1}, e_{i+1}, \ldots, e_{N-1} \leftarrow \chi_{\sigma_1} \\ z_0, \ldots, z_{i-2} \leftarrow \mathcal{U}(R_q), z_{i-1} = z_{i+1} - r_i, z_i \leftarrow \mathcal{U}(R_q), \\ z_{i+1} = a s_{i+1} + e_{i+1}, \ldots, z_{N-1} = a s_{N-1} + e_{N-1}; \\ e_1', \ldots, e_{N-1}' \leftarrow \chi_{\sigma_1}; \hat{e}_0 \leftarrow \chi_{\sigma_2} \\ X_0' = -\sum_{i=1}^{N-1} X_i + \hat{e}_0, X_1, \ldots, X_{i-1} \leftarrow \mathcal{U}(R_q), \qquad : (\mathsf{T}, \mathsf{sk}) \\ X_i \leftarrow \mathcal{U}(R_q), X_{i+1} = (z_{i+2} - z_i) s_{i+1} + e_{i+1}', \ldots, \\ \quad X_{N-1} = (z_0 - z_{N-2}) s_{N-1} + e_{N-1}'; e_{N-1}'' \leftarrow \chi_{\sigma_1}; \\ b_{N-1} = z_{N-2} N s_{N-1} + e_{N-1}'' + X_{N-1} \cdot (N-1) + \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\ (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}). \end{array} \right\}$$

Experiment $3i+2$. This experiment proceeds exactly the same as *Experiment* $3i+1$, except that z_{i-1} is directly sampled from $\mathcal{U}(R_q)$. The corresponding distribution is denoted as Dist_{3i+2}. We leave the formal definition of Dist_{3i+2} implicit for simplicity.

Bounding the difference of $|\Pr_{3i}[\mathsf{Query}]-\Pr_{3i-1}[\mathsf{Query}]|$, $|\Pr_{3i+1}[\mathsf{Query}]-\Pr_{3i}[\mathsf{Query}]|$, and $|\Pr_{3i+2}[\mathsf{Query}]-\Pr_{3i+1}[\mathsf{Query}]|$ follows exactly the same logic as *bounding the differences of* $|\Pr_3[\mathsf{Query}]-\Pr_2[\mathsf{Query}]|$, $|\Pr_4[\mathsf{Query}]-\Pr_3[\mathsf{Query}]|$, and $|\Pr_5[\mathsf{Query}]-\Pr_4[\mathsf{Query}]|$, respectively. Then we have

$$\Pr_{3i}[\mathsf{Query}] = \Pr_{3i-1}[\mathsf{Query}]; \tag{17}$$

$$|\Pr_{3i+1}[\mathsf{Query}] - \Pr_{3i}[\mathsf{Query}]| \le \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1); \tag{18}$$

$$\Pr_{3i+2}[\mathsf{Query}] = \Pr_{3i+1}[\mathsf{Query}]; \tag{19}$$

Note that in *Experiment* $3N-4$, the last experiment of the experiment sequence above, we already switched all the z_i, X_i up to z_{N-1}, X_{N-1}. We construct the next two experiments to switch $z_{N-1}, X_{N-1}, b_{N-1}$.

Experiment $3N-3$**.** The experiment proceeds exactly the same as *Experiment* $3N-4$, except that we let $z_{N-2}=r_2, X_{N-1}=r_1 s_{N-1}+e'_{N-1}, z_0=r_1+r_2$, where r_1, r_2 are uniformly distributed in R_q. The corresponding distribution is as follows, denoted Dist_{3N-3}.

Bounding the difference of $|\Pr_{3N-3}[\mathsf{Query}]-\Pr_{3N-4}[\mathsf{Query}]|$:

Since r_1, r_2 is sampled uniformly, r_1+r_2 is also uniformly distributed in R_q. Then we claim that *Experiment* $3N-3$ is identical to *Experiment* $3N-4$ up to variable substitution, written as

$$\Pr_{3N-3}[\mathsf{Query}] = \Pr_{3N-4}[\mathsf{Query}]; \tag{20}$$

$$\mathsf{Dist}_{3N-3} := \left\{ \begin{array}{l} a, r_1, r_2 \leftarrow \mathcal{U}(R_q), s_{N-1}, e_{N-1} \leftarrow \chi_{\sigma_1}; z_0 = r_1 + r_2, \\ z_1, \ldots, z_{N-3} \leftarrow \mathcal{U}(R_q), z_{N-2} = r_2, \\ z_{N-1} = a s_{N-1} + e_{N-1}; \hat{e}_0 \leftarrow \chi_{\sigma_2}; e'_{N-1} \leftarrow \chi_{\sigma_1}; \\ X'_0 = -\sum_{i=1}^{N-1} X_i + \hat{e}_0, X_1, \ldots, X_{N-2} \leftarrow \mathcal{U}(R_q), \\ X_{N-1} = r_1 s_{N-1} + e'_{N-1}; e''_{N-1} \leftarrow \chi_{\sigma_1}; \\ b_{N-1} = r_2 N s_{N-1} + e''_{N-1} + X_{N-1} \cdot (N-1) + \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\ (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}). \end{array} : (\mathsf{T}, \mathsf{sk}) \right\}$$

Experiment $3N-2$**.** This experiment proceeds exactly the same as *Experiment* $3N-3$, except that $z_{N-1}, X_{N-1}, b_{N-1}$ are generated from $\mathcal{U}(R_q)$. The corresponding distribution is as follows, denoted Dist_{3N-2}:

$$\mathsf{Dist}_{3N-2} := \left\{ \begin{array}{l} a \leftarrow \mathcal{U}(R_q), z_0, z_1 \ldots, z_{N-2} \leftarrow \mathcal{U}(R_q), \\ z_{N-1} \leftarrow \mathcal{U}(R_q); \hat{e}_0 \leftarrow \chi_{\sigma_2}; r_1, r_2 \leftarrow \mathcal{U}(R_q) \\ X_0' = -\sum_{i=1}^{N-1} X_i + \hat{e}_0, X_1, \ldots, X_{N-1} \leftarrow \mathcal{U}(R_q) \qquad : (\mathsf{T}, \mathsf{sk}) \\ b_{N-1} \leftarrow \mathcal{U}(R_q); \\ (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}). \end{array} \right\}$$

Bounding the difference of $|\Pr_{3N-2}[\mathsf{Query}] - \Pr_{3N-3}[\mathsf{Query}]|$:

Let $b_{rlwe} = r_2 N s_{N-1} + e''_{N-1}$, then $b_{N-1} = b_{rlwe} + X_{N-1} \cdot (N-1) + X_0 \cdot (N-2) + \cdots + X_{N-3}$. As r_2 is sampled uniformly at random and N is invertible over R_q, $r_2 N$ is uniformly distributed in R_q.

Given an algorithm $\mathcal{A}$ running in time t attacking group key exchange protocol $\mathit{\Pi}$, let $\mathcal{B}$ be an algorithm that takes as input (a, z_{N-1}), (r_1, X_{N-1}), and $(r_2 N, b_{rlwe})$, generates $(\mathsf{T}, \mathsf{sk})$ based on distribution Dist'_{3N-3} which is identical to Dist_{3N-3} except for (a, z_{N-1}), (r_1, X_{N-1}), and $(r_2 N, b_{rlwe})$ given as input. $\mathcal{B}$ runs $\mathcal{A}$ as subroutine and outputs whatever $\mathcal{A}$ outputs. Note that running time t_1 of $\mathcal{B}$ equals to t plus a minor overhead for the simulation of the security experiment for $\mathcal{A}$.

It is straightforward to see that if (a, z_{N-1}), (r_1, X_1), and $(r_2 N, b_{rlwe})$ are sampled from the Ring-LWE distribution $A_{n,q,\chi_{\sigma_1}}$, then Dist'_{3N-3} is identical to Dist_{3N-3}. If (a, z_{N-1}), (r_1, X_{N-1}), and $(r_2 N, b_{rlwe})$ are sampled from $\mathcal{U}(R_q^2)$, then Dist'_{3N-3} is identical to Dist_{3N-2}, since when b_{rlwe} is sampled uniformly at random, $b_{rlwe} + X_{N-1} \cdot (N-1) + X_0 \cdot (N-2) + \cdots + X_{N-3}$ is also uniformly distributed over R_q.

Therefore we conclude that the difference of algorithm $\mathcal{A}^{\mathsf{GKE}}$'s success probability in *Experiment 3N - 2* and *Experiment 3N - 3* is bounded by the advantage of adversary $\mathcal{B}$ running in time t_1 in distinguishing Ring-LWE from $\mathcal{U}(R_q)$ given three samples. Thus, we conclude that

$$|\Pr_{3N-2}[\mathsf{Query}] - \Pr_{3N-3}[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1). \tag{21}$$

Experiment $3N-1$. This experiment proceeds exactly the same as *Experiment* $3N-2$, except that k_{N-1} is directly sampled uniformly from $\{0,1\}^\lambda$. Note that the corresponding distribution is exactly the distribution Ideal.

$$\mathsf{Ideal} := \left\{ \begin{array}{l} a \leftarrow \mathcal{U}(R_q); z_0, \ldots, z_{N-1} \leftarrow \mathcal{U}(R_q); e_0' \leftarrow \chi_{\sigma_1}; \\ X_0' = -\sum_{i=1}^{N-1} X_i + \hat{e}_0, X_1, \ldots, X_{N-1} \leftarrow \mathcal{U}(R_q) \\ b_{N-1} \leftarrow \mathcal{U}(R_q); (K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}) \qquad : (\mathsf{T}, \mathsf{sk}) \\ k'_{N-1} \leftarrow \{0,1\}^\lambda; \mathsf{sk} = \mathcal{H}(k'_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0', \ldots, X_{N-1}, K_{N-1}); \end{array} \right\}$$

Bounding the difference of $|\Pr_{3N-1}[\mathsf{Query}] - \Pr_{3N-2}[\mathsf{Query}]|$:

Given transcript T, and b_{N-1} which is uniformly distributed, using a straight forward reduction, we obtain advantage of adversary $\mathcal{B}$ running in time t_2 in distinguishing k_{N-1} computed by $\mathsf{recMsg}(b_{N-1})$ from a uniform bit string k'_{N-1} with length λ is at least $|\Pr_{3N-1}[\mathsf{Query}] - \Pr_{3N-2}[\mathsf{Query}]|$, namely,

$$|\Pr_{3N-1}[\mathsf{Query}] - \Pr_{3N-2}[\mathsf{Query}]| \leq \mathsf{Adv}_{\mathsf{KeyRec}}(t_2). \tag{22}$$

Note that t_2 equals to the running time of adversary $\mathcal{A}$ attacking the protocol Π, plus a minor overhead for simulating experiment for $\mathcal{A}$.

Finally, since adversary attacking the GKE protocol Π makes at most q queries to the random oracle, $\Pr_{3N-1}[\mathsf{Query}] = \frac{\mathsf{q}}{2^\lambda} \in \mathsf{negl}(\lambda)$. Combining Equations (13) - (22), we have

$$\Pr_1[\mathsf{Query}] \leq N \cdot \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1) + \mathsf{Adv}_{\mathsf{KeyRec}}(t_2) + \frac{\mathsf{q}}{2^\lambda}. \tag{23}$$

The theorem now follows immediately from Equations (1), (2), and (23). $\square$